

APPENDIX B

Appendix listing of additional Click modules ("elements").

ADAPTIVESHAPER(n)

ADAPTIVESHAPER(n)

NAME

AdaptiveShaper - Click element

SYNOPSIS

AdaptiveShaper(DROP_P, REPRESS_WEIGHT)

PROCESSING TYPE

Push

DESCRIPTION

AdaptiveShaper is a push element that shapes input traffic from input port 0 to output port 0. Packets are shaped based on "repressive" traffic from input port 1 to output port 1. Each repressive packet increases a multiplicative factor f by `REPRESS_WEIGHT`. Each input packet is killed instead of pushed out with $f * \text{DROP_P}$ probability. After each dropped packet, f is decremented by 1.

EXAMPLES

ELEMENT HANDLERS

`drop_prob` (read/write)
value of `DROP_P`

`repress_weight` (read/write)
value of `REPRESS_WEIGHT`

SEE ALSO

`PacketShaper(n)`, `RatioShaper(n)`

FOIA b 7 - DATED 03-16-2001

APPENDIX B

ADAPTIVESPLITTER(n)

ADAPTIVESPLITTER(n)

NAME

AdaptiveSplitter - Click element

SYNOPSIS

AdaptiveSplitter(RATE)

PROCESSING TYPE

Push

DESCRIPTION

AdaptiveSplitter attempts to split RATE number of packets per second for each address. It takes the fwd_rate annotation set by IPRateMonitor(n), and calculates a split probability based on that rate. The split probability attempts to guarantee RATE number of packets per second. That is, the lower the fwd_rate, the higher the split probability.

Splitted packets are on output port 1. Other packets are on output port 0.

EXAMPLES

AdaptiveSplitter(10);

SEE ALSO

IPRateMonitor(n)

00931291-001601
T09T20-T62T660

APPENDIX B

ADDRFILTER(n)

ADDRFILTER(n)

NAME

AddrFilter - Click element

SYNOPSIS

AddrFilter(DST/SRC, N)

PROCESSING TYPE

Push

DESCRIPTION

Filters out IP addresses given in write handler. DST/SRC specifies which IP address (dst or src) to filter. N is the maximum number of IP addresses to filter at any time. Packets passed the filter goes to output 0. Packets rejected by the filter goes to output 1.

AddrFilter looks at addresses in the IP header of the packet, not the annotation. It requires an IP header annotation (MarkIPHeader(n)).

EXAMPLES

AddrFilter(DST, 8)

Filters by dst IP address, up to 8 addresses.

ELEMENT HANDLERS

table ((read))

Dumps the list of addresses to filter and

add ((write))

Expects a string "addr mask duration", where addr is an IP address, mask is a netmask, and duration is the number of seconds to filter packets from this IP address. If 0 is given as a duration, filtering is removed. For example, "18.26.4.0 255.255.255.0 10" would filter out all packets with dst or source address 18.26.4.* for 10 seconds. New addresses push out old addresses if more than N number of filters already exist.

reset ((write))

Resets on write.

SEE ALSO

Classifier(n), MarkIPHeader(n)

00931291-081601
T09T80" T62TE660

APPENDIX B

ATTACKLOG (n)

ATTACKLOG (n)

NAME

AttackLog - Click element; maintains a log of attack packets in SAVE_FILE.

SYNOPSIS

AttackLog(SAVE_FILE, INDEX_FILE, MULTIPLIER, PERIOD)

PROCESSING TYPE

Agnostic

DESCRIPTION

Maintains a log of attack packets in SAVE_FILE. Expects packets with ethernet headers, but with the first byte of the ethernet header replaced by an attack bitmap, set in kernel. AttackLog classifies each packet by the type of attack, and maintains an attack rate for each type of attack. The attack rate is the arrival rate of attack packets multiplied by MULTIPLIER.

AttackLog writes a block of data into SAVE_FILE once every PERIOD number of seconds. Each block is composed of entries of the following format:

delimiter (0s)	4 bytes
time	4 bytes
attack type	2 bytes
attack rate	4 bytes
ip header and payload (padded)	86 bytes

	100 bytes

Entries with the same attack type are written out together. A delimiter of 0xFFFFFFFF is written to the end of each block.

A circular timed index file is kept in INDEX_FILE along side the attacklog. See CircularIndex(n).

SEE ALSO

CircularIndex(n)

09931291-001601
T09T00-T06T00

APPENDIX B

CIRCULARINDEX (n)

CIRCULARINDEX (n)

NAME

CircularIndex - Click element; writes a timed circular index into a file.

SYNOPSIS

CircularIndex

DESCRIPTION

CircularIndex writes an entry into a circular index file periodically. The entry contains a 32 bit time stamp and a 64 bit offset into another file. The following functions are exported by CircularIndex.

int initialize(String FILE, unsigned PERIOD, unsigned WRAP) - Use FILE as the name of the circular file. Writes entry into circular file once every PERIOD number of seconds. WRAP is the number of writes before wrap around. If WRAP is 0, the file is never wrapped around.

void write_entry(long long offset) - Write entry into index file. Use offset as the offset in the entry.

SEE ALSO

GatherRates(n), MonitorSRC16(n)

FOUO "T62FE60"

APPENDIX B

DISCARDTODEVICE (n)

DISCARDTODEVICE (n)

NAME

DiscardToDevice - Click element; drops all packets. gives
skbs to device.

SYNOPSIS

DiscardToDevice (DEVICE)

PROCESSING TYPE

Agnostic

DESCRIPTION

Discards all packets received on its single input. Gives
all skbuffs to specified device.

09031291.081601
T09T80" T62TE660

APPENDIX B

FILTERTCP (n)

FILTERTCP (n)

NAME

FilterTCP - Click element

SYNOPSIS

FilterTCP()

PROCESSING TYPE

Push

DESCRIPTION

Expects TCP/IP packets as input.

09931201-081601
T09T80" T62T660

APPENDIX B

FROMTUNNEL (n)

FROMTUNNEL (n)

NAME

FromTunnel - Click element

SYNOPSIS

FromTunnel(TUNNEL, SIZE, BURST)

PROCESSING TYPE

Push

DESCRIPTION

Grab packets from kernel KUTunnel element. TUNNEL is a /proc file in the handler directory of the KUTunnel element. SIZE specifies size of the buffer to use (if packet in kernel has larger size, it is dropped). BURST specifies the maximum number of packets to push each time FromTunnel runs.

EXAMPLES

FromTunnel(/proc/click/tunnel/config)

09031291.081601
T09T00" T62TE660

APPENDIX B

GATHERRATES (n)

GATHERRATES (n)

NAME

GatherRates - Click element

SYNOPSIS

```
GatherRates(SAVE_FILE, INDEX_FILE, TCPMONITOR_IN, TCPMONI-
TOR_OUT, MONITOR_PERIOD, SAVE_PERIOD);
```

PROCESSING TYPE

Agnostic

DESCRIPTION

Gathers aggregate traffic rates from TCPMonitor(n) element at TCPMONITOR_IN and TCPMONITOR_OUT.

Aggregate rates are gathered once every MONITOR_PERIOD number of seconds. They are averaged and saved to SAVE_FILE once every SAVE_PERIOD number of seconds. The following entry is written to SAVE_FILE for both incoming and outgoing traffic:

delimiter (0s)	4 bytes
time	4 bytes
type (0 for incoming traffic, 1 for outgoing traffic)	4 bytes
packet rate of tcp traffic	4 bytes
byte rate of tcp traffic	4 bytes
rate of fragmented tcp packets	4 bytes
rate of tcp syn packets	4 bytes
rate of tcp fin packets	4 bytes
rate of tcp ack packets	4 bytes
rate of tcp rst packets	4 bytes
rate of tcp psh packets	4 bytes
rate of tcp urg packets	4 bytes
packet rate of non-tcp traffic	4 bytes
byte rate of non-tcp traffic	4 bytes
rate of fragmented non-tcp traffic	4 bytes
rate of udp packets	4 bytes
rate of icmp packets	4 bytes
rate of all other packets	4 bytes

	72 bytes

After the two entries, an additional delimiter of 0xFFFFFFFF is written. SAVE_PERIOD must be a multiple of MONITOR_PERIOD.

A circular timed index is kept along side the stats file. See CircularIndex(n).

SEE ALSO

TCPMonitor(n) CircularIndex(n)

APPENDIX B

ICMPPINGENCAP (n)

ICMPPINGENCAP (n)

NAME

ICMPPINGEncap - Click element

SYNOPSIS

ICMPPINGEncap(SADDR, DADDR [, CHECKSUM?])

DESCRIPTION

Encapsulates each incoming packet in a ICMP ECHO/IP packet with source address SADDR and destination address DADDR. The ICMP and IP checksums are calculated if CHECKSUM? is true; it is true by default.

EXAMPLES

ICMPPINGEncap(1.0.0.1, 2.0.0.2)

09931291-081601
T09T80-T62T660

APPENDIX B

KUTUNNEL(n)

KUTUNNEL(n)

NAME

KUTunnel - Click element; stores packets in a FIFO queue that userlevel Click elements pull from.

SYNOPSIS

KUTunnel([CAPACITY])

PROCESSING TYPE

Push

DESCRIPTION

Stores incoming packets in a first-in-first-out queue. Drops incoming packets if the queue already holds CAPACITY packets. The default for CAPACITY is 1000. Allows user-level elements to pull from queue via ioctl.

ELEMENT HANDLERS

length (read-only)

Returns the current number of packets in the queue.

highwater_length (read-only)

Returns the maximum number of packets that have ever been in the queue at once.

capacity (read/write)

Returns or sets the queue's capacity.

drops (read-only)

Returns the number of packets dropped so far.

SEE ALSO

Queue(n)

B-12

LOGGER (n)

Logger - Click element

```
Logger(LOGFILE, INDEXFILE [, LOCKFILE, COMPRESS?, LOGSIZE,
PACKETSIZE, WRITEPERIOD, IDXCOALESC, PACKETFREQ, MAXBUF-
SIZE ] )
```

Agnostic

Has one input and one output.

```

-----
|   time (6 bytes)   |
|  length (2 bytes)  |
|   packet data      |

```

INDEXFILE maintains control data for LOGFILE. It contains a sequence of sequence control blocks of the following form:

```

-----
|   date (4 bytes)   |
| offset (sizeof off_t) |
| length (sizeof off_t) |

```

Sequence control blocks are always stored in increasing chronological order; offsets need not be in increasing order, since LOGFILE is a circular buffer.

COMPRESS? (true, false) determines whether packet data is logged in compressed form. Default is true.

APPENDIX B

LOGSIZE specifies the maximum allowable log file size, in KB. Default is 2GB. LOGSIZE=0 means "grow as necessary".

PACKETSIZE is the amount of packet data stored in the log. By default, the first 120 (128-6-2) bytes are logged and the remainder is discarded. Note that PACKETSIZE is the amount of data logged before compression.

Packet records are buffered in memory and periodically written to LOGFILE as a packet sequence. WRITEPERIOD is the number of seconds that should elapse between writes to LOGFILE. Default is 60. INDEXFILE is updated every time a sequence of buffered packet records is written to LOGFILE. The date in the sequence control block is the time of the first packet record of the sequence, with milliseconds omitted.

PACKETFREQ is an estimate of the number of packets per second that will be passing through Logger. Combined with WRITEPERIOD, this is a hint of buffer memory requirements. By default, PACKETFREQ is 1000. Since by default WRITEPERIOD is 60 and each packet record is at most 128 bytes, Logger normally allocates 7500KB of memory for the buffer. Logger will grow the memory buffer as needed up to a maximum of MAXBUFSIZE KB, at which point the buffered packet records are written to disk even if WRITEPERIOD seconds have not elapsed since the last write. Default MAXBUFSIZE is 65536 (64MB).

T09T80" T62T660

APPENDIX B

MONITORSRC16(n)

MONITORSRC16(n)

NAME

MonitorSRC16 - Click element

SYNOPSIS

```
MonitorSRC16(SAVE_FILE, INDEX_FILE, MULTIPLIER, PERIOD,
             WRAP)
```

PROCESSING TYPE

Agnostic

DESCRIPTION

Examines src address of packets passing by. Collects statistics for each 16 bit IP address prefix. The following data structure is written to SAVE_FILE for every 16 bit IP address prefix every PERIOD number of seconds.

delimiter (0s)	(4 bytes)
time	(4 bytes)
addr	(4 bytes)
tcp rate	(4 bytes)
non tcp rate	(4 bytes)
percent of tcp	(1 byte)
percent of tcp frag	(1 byte)
percent of tcp syn	(1 byte)
percent of tcp fin	(1 byte)
percent of tcp ack	(1 byte)
percent of tcp rst	(1 byte)
percent of tcp psh	(1 byte)
percent of tcp urg	(1 byte)
percent of non tcp frag	(1 byte)
percent of udp	(1 byte)
percent of icmp	(1 byte)
reserved	(1 byte)

32 bytes

TCP and non TCP rates are multiplied by MULTIPLIER. An additional delimiter of 0xFFFFFFFF is written at the end of a block of entries.

WARP specifies the number of writes before wrap-around. For example, if PERIOD is 60, WARP is 5, then every 5 minutes, the stats file wrap around.

A timed circular index is maintained along side the statistics file in INDEX_FILE. See CircularIndex(n).

SEE ALSO

CircularIndex(n)

APPENDIX B

RANDOMTCPIPENCAP(n)

RANDOMTCPIPENCAP(n)

NAME

RandomTCPIPEncap - Click element

SYNOPSIS

RandomTCPIPEncap(DA BITS [DP SEQN ACKN CHECKSUM SA MASK])

PROCESSING TYPE

Agnostic

DESCRIPTION

Encapsulates each incoming packet in a TCP/IP packet with random source address and source port, destination address DA, and control bits BITS. If BITS is -1, control bits are also generated randomly. If destination port DP, sequence number SEQN, or ack number ACKN is specified and non-zero, it is used. Otherwise, it is generated randomly for each packet. IP and TCP checksums are calculated if CHECKSUM is true; it is true by default. SEQN and ACKN should be in host order. SA and MASK are optional IP address; if they are specified, the source address is computed as ((random() & MASK) | SA).

EXAMPLES

RandomTCPIPEncap(1.0.0.2 4)

SEE ALSO

RoundRobinTCPIPEncap(n), RandomUDPIPEncap(n)

09031291-081601
T09T80-T62TE660

APPENDIX B

RANDOMUDPIPCAP (n)

RANDOMUDPIPCAP (n)

NAME

RandomUDPIPCap - Click element

SYNOPSIS

RandomUDPIPCap(SADDR SPORT DADDR DPORT PROB [CHECKSUM?]
[, ...])

PROCESSING TYPE

Agnostic

DESCRIPTION

Encapsulates each incoming packet in a UDP/IP packet with source address SADDR, source port SPORT, destination address DADDR, and destination port DPORT. The UDP checksum is calculated if CHECKSUM? is true; it is true by default.

PROB gives the relative chance of this argument be used over others.

The RandomUDPIPCap element adds both a UDP header and an IP header.

You can a maximum of 16 arguments. Each argument specifies a single UDP/IP header. The element will randomly pick one argument. The relative probabilities are determined by PROB.

The Strip(n) element can be used by the receiver to get rid of the encapsulation header.

EXAMPLES

```
RandomUDPIPCap(1.0.0.1 1234 2.0.0.2 1234 1 1,  
               1.0.0.2 1093 2.0.0.2 1234 2 1)
```

Will send about twice as much UDP/IP packets with 1.0.0.2 as its source address than packets with 1.0.0.1 as its source address.

SEE ALSO

Strip(n), UDPIPCap(n), RoundRobinUDPIPCap(n)

10931291-031601

APPENDIX B

RATEWARN (n)

RATEWARN (n)

NAME

RateWarn - Click element; classifies traffic and sends out warnings when rate of traffic exceeds specified rate.

SYNOPSIS

RateWarn(RATE, WARNFREQ)

PROCESSING TYPE

Push

DESCRIPTION

RateWarn has three output ports. It monitors the rate of packet arrival on input port 0. Packets are forwarded to output port 0 if rate is below RATE. If rate exceeds RATE, it sends out a warning packet WARNFREQ number of seconds apart on output port 2 in addition to forwarding all traffic through output port 1.

SEE ALSO

PacketMeter(n)

09931291-081601
T09T80-T62T660

APPENDIX B

RATIOSHAPER(n)

RATIOSHAPER(n)

NAME

RatioShaper - Click element

SYNOPSIS

RatioShaper(FWD_WEIGHT, REV_WEIGHT, THRESH, P)

PROCESSING TYPE

Push

DESCRIPTION

RatioShaper shapes packets based on fwd_rate_anno and rev_rate_anno rate annotations set by IPRateMonitor(n). If either annotation is greater than THRESH, and $FWD_WEIGHT * fwd_rate_anno > REV_WEIGHT * rev_rate_anno$, the packet is moved onto output port 1 with a probability of

$$\min(1, P * (fwd_rate_anno * FWD_WEIGHT) / (rev_rate_anno * REV_WEIGHT))$$

FWD_WEIGHT, REV_WEIGHT, and THRESH are integers. P is a decimal between 0 and 1. Otherwise, packet is forwarded on output port 0.

EXAMPLES

RatioShaper(1, 2, 100, .2);

if fwd_rate_anno more than twice as big as rev_rate_anno, and both rates are above 100, drop packets with an initial probability of 20 percent.

ELEMENT HANDLERS

fwd_weight (read/write)
value of FWD_WEIGHT

rev_weight (read/write)
value of REV_WEIGHT

thresh (read/write)
value of THRESH

drop_prob (read/write)
value of P

SEE ALSO

Block(n), IPRateMonitor(n)

APPENDIX B

REPORTACTIVITY (n)

REPORTACTIVITY (n)

NAME

ReportActivity - Click element

SYNOPSIS

ReportActivity(SAVE_FILE, IDLE)

PROCESSING TYPE

Agnostic

DESCRIPTION

Write into SAVE_FILE a 32 bit time value followed by an ASCII representation of that time stamp whenever a packet comes by. If IDLE number of seconds pass by w/o a packet, removes the file.

09931291.081601

APPENDIX B

ROUNDROBINSETIPADDRESS (n)

ROUNDROBINSETIPADDRESS (n)

NAME

RoundRobinSetIPAddress - Click element

SYNOPSIS

RoundRobinSetIPAddress(ADDR [, ...])

PROCESSING TYPE

Agnostic

DESCRIPTION

Set the destination IP address annotation of each packet with an address chosen from the configuration string in round robin fashion. Does not compute checksum (use SetIPChecksum(n) or SetUDPTCPChecksum(n)) or encapsulate the packet with headers (use RoundRobinUDPIPEncap(n) or RoundRobinTCPIPEncap(n) with bogus address).

EXAMPLES

```
RoundRobinUDPIPEncap(2.0.0.2 0.0.0.0 0 0 0)
-> RoundRobinSetIPAddress(1.0.0.2, 1.0.0.3, 1.0.0.4)
-> StoreIPAddress(12)
-> SetIPChecksum
-> SetUDPTCPChecksum
```

this configuration segment places an UDP header onto each packet, with randomly generated source and destination ports. The destination IP address is 2.0.0.2, the source IP address is 1.0.0.2, or 1.0.0.3, or 1.0.0.4. Both IP and UDP checksum are computed.

SEE ALSO

RoundRobinUDPIPEncap(n), RoundRobinTCPIPEncap(n), UDPIPEncap(n), SetIPChecksum(n), SetUDPTCPChecksum(n), SetIPAddress(n), StoreIPAddress(n)

09931291.031601

APPENDIX B

ROUNDROBINTCPIPEncap(n)

ROUNDROBINTCPIPEncap(n)

NAME

RoundRobinTCPIPEncap - Click element

SYNOPSIS

RoundRobinTCPIPEncap(SA DA BITS [SP DP SEQN ACKN CHECKSUM]
[, ...])

PROCESSING TYPE

Agnostic

DESCRIPTION

Encapsulates each incoming packet in a TCP/IP packet with source address SA, source port SP (if 0, a random one is generated for each packet), destination address DA, and destination port DP (if 0, a random one is generated for each packet), and control bits BITS. If SEQN and ACKN specified are non-zero, they are used. Otherwise, they are randomly generated for each packet. IP and TCP checksums are calculated if CHECKSUM is true; it is true by default. SEQN and ACKN should be in host order.

The RoundRobinTCPIPEncap element adds both a TCP header and an IP header.

You can give as many arguments as you'd like. Each argument specifies a single TCP/IP header. The element will cycle through the headers in round-robin order.

The Strip(n) element can be used by the receiver to get rid of the encapsulation header.

EXAMPLES

RoundRobinTCPIPEncap(2.0.0.2 1.0.0.2 4 1022 1234 42387492 2394839 1,
2.0.0.2 1.0.0.2 2)

SEE ALSO

Strip(n), RoundRobinUDPIPEncap(n)

0931291-031601

APPENDIX B

ROUNDROBINUDPIPECAP (n)

ROUNDROBINUDPIPECAP (n)

NAME

RoundRobinUDPIPEncap - Click element

SYNOPSIS

RoundRobinUDPIPEncap(SADDR DADDR [SPORT DPORT CHECKSUM?]
[, ...])

PROCESSING TYPE

Agnostic

DESCRIPTION

Encapsulates each incoming packet in a UDP/IP packet with source address SADDR, source port SPORT, destination address DADDR, and destination port DPORT. The UDP and IP checksums are calculated if CHECKSUM? is true; it is true by default. If either DPORT or SPORT is 0, the port will be randomly generated for each packet.

The RoundRobinUDPIPEncap element adds both a UDP header and an IP header.

You can give as many arguments as you'd like. Each argument specifies a single UDP/IP header. The element will cycle through the headers in round-robin order.

The Strip(n) element can be used by the receiver to get rid of the encapsulation header.

EXAMPLES

RoundRobinUDPIPEncap(2.0.0.2 1.0.0.2 1234 1002 1,
2.0.0.2 1.0.0.2 1234)

SEE ALSO

Strip(n), UDPIPEncap(n)

09931291-081601
TOTAL 081601

APPENDIX B

SETSNIFFFLAGS (n)

SETSNIFFFLAGS (n)

NAME

SetSniffFlags - Click element; sets sniff flags annotation.

SYNOPSIS

SetSniffFlags(FLAGS [, CLEAR])

PROCESSING TYPE

Agnostic

DESCRIPTION

Set the sniff flags annotation of incoming packets to FLAGS bitwise or with the old flags. if CLEAR is true (false by default), the old flags are ignored.

09931591.081501
T09T80.T62E60

APPENDIX B

SETUDPTCPChecksum(n)

SETUDPTCPChecksum(n)

NAME

SetUDPTCPChecksum - Click element

SYNOPSIS

SetUDPTCPChecksum()

PROCESSING TYPE

Agnostic

DESCRIPTION

Expects an IP packet as input. Calculates the ICMP, UDP or TCP header's checksum and sets the checksum header field. Does not modify packet if it is not an ICMP, UDP, or TCP packet.

SEE ALSO

SetIPChecksum(n)

0931291-081601
T09T80-T62TE60

APPENDIX B

STORESNIFFFLAGS (n)

STORESNIFFFLAGS (n)

NAME

StoreSniffFlags - Click element; stores sniff flags annotation in packet

SYNOPSIS

StoreSniffFlags(OFFSET)

PROCESSING TYPE

Agnostic

DESCRIPTION

Copy the sniff flags annotation into the packet at offset OFFSET.

09931291 "081601

APPENDIX B

TCPMONITOR(n)

TCPMONITOR(n)

NAME

TCPMonitor - Click element

SYNOPSIS

TCPMonitor()

PROCESSING TYPE

Push

DESCRIPTION

Monitors and splits TCP traffic. Output 0 are TCP traffic, output 1 are non-TCP traffic. Keeps rates of TCP, TCP BYTE, SYN, ACK, PUSH, RST, FIN, URG, and fragmented packets. Also keeps rates of ICMP, UDP, non-TCP BYTE, and non-TCP fragmented traffic.

ELEMENT HANDLERS

rates (read)
dumps rates

09031291-081601
T09T30-T62E650

APPENDIX B

TCPSYNPROXY (n)

TCPSYNPROXY (n)

NAME

TCPSYNProxy - Click element

SYNOPSIS

```
TCPSYNProxy(MAX_CONNS, THRESHOLD, MIN_TIMEOUT, MAX_TIMEOUT
[, PASSIVE])
```

PROCESSING TYPE

Push

DESCRIPTION

Help setup a three way TCP handshake from A to B by supplying the last ACK packet to the SYN ACK B sent prematurely, and send RST packets to B later if no ACK was received from A.

Expects IP encapsulated TCP packets, each with its ip header marked (MarkIPHeader(n) or CheckIPHeader(n)).

Aside from responding to SYN ACK packets from B, TCPSYNProxy also examines SYN packets from A. When a SYN packet from A is received, if there are more than MAX_CONNS number of outstanding 3 way connections per destination (daddr + dport), reject the SYN packet. If MAX_CONNS is 0, no maximum is set.

The duration from sending an ACK packet to B to sending a RST packet to B decreases exponentially as the number of outstanding connections to B increases pass $2^{\text{THRESHOLD}}$. The minimum timeout is MIN_TIMEOUT. If the number of outstanding half-open connections is above $2^{\text{THRESHOLD}}$, the timeout is

$$\max(\text{MIN_TIMEOUT}, \text{MAX_TIMEOUT} \gg (N \gg \text{THRESHOLD}))$$

Where N is the number of outstanding half-open connections. For example, let the MIN_TIMEOUT value be 4 seconds, the MAX_TIMEOUT value be 90 seconds, and THRESHOLD be 3. Then when $N < 8$, timeout is 90. When $N < 16$, timeout is 45. When $N < 24$, timeout is 22 seconds. When $N < 32$, timeout is 11 seconds. When $N < 64$, timeout is 4 seconds. Timeout period does not decrement if the threshold is 0.

TCPSYNProxy has two inputs, three outputs. All inputs and outputs take in and spew out packets with IP header. Input 0 expects TCP packets from A to B. Input 1 expects TCP packets from B to A. Output 0 spews out packets from A to B. Output 1 spews out packets from B to A. Output 2 spews out the ACK and RST packets generated by the element.

If PASSIVE is true (it is not by default), monitor TCP three-way handshake instead of actively setting it up. In

09931291-081601

APPENDIX B

this case, no ACK or RST packets will be sent. When an outstanding SYN times out, the SYN ACK packet is sent out of output port 2. No packets on port 0 are modified or dropped in this operating mode.

EXAMPLES

... -> CheckIPHeader() -> TCPSYNProxy(128,3,10,90) -> ...

ELEMENT HANDLERS

summary (read)

Returns number of ACK and RST packets sent and number of SYN packets rejected.

table (read)

Dumps the table of half-opened connections.

reset (write)

Resets on write.

SEE ALSO

MarkIPHeader(n), CheckIPHeader(n)

09931291-081601
109180-1627E660

APPENDIX B

TCPSYNRESP (n)

TCPSYNRESP (n)

NAME

TCPSYNResp - Click element

SYNOPSIS

TCPSYNResp()

PROCESSING TYPE

Push

DESCRIPTION

Takes in TCP packet, if it is a SYN packet, return a SYN ACK. This is solely for debugging and performance tuning purposes. No checksum is done. Spews out original packet on output 0 untouched. Spews out new packet on output 1.

201094509.doc

09931291-031601
T09T80-T62T660